QUANTUM COMPUTING: REALITY OR HYPE?

Neal Koblitz, University of Washington, Seattle

Quantum computing was described as a theoretical possibility about 40 years ago in writings by three physicists (including the Nobel Prize winner Richard Feynman) and one mathematician, Yuri Manin. The idea started to attract more interest in 1994, when Peter Shor developed an algorithm for a quantum computer that would factor large integers rapidly and find logarithms in finite systems. This would break both of the most widely used types of public-key cryptography — RSA and Elliptic Curve Cryptography (ECC).

Because of the possibility that someone succeeds in building a quantum computer on a scale that can defeat RSA and ECC cryptography, a lot of work in recent years has been devoted to developing "quantum-safe" methods of encryption, digital signatures, and key exchange that resist quantum attacks. The U.S. government agency NIST (National Institute for Standards and Technology) has been conducting a competition for the best quantum-safe cryptographic systems. NIST plans to announce the winning cryptosystems later this month (March 2022).

Starting in the late 1990s several academic and industrial research centers have devoted many millions of dollars to efforts to develop quantum computing. But progress has been very, very slow.

Formidable Obstacles to Construction

According to current recommendations for key sizes, the numbers N used in RSA should have 2048 bits — that is, over 600 decimal digits. Such numbers are too large to be factored by today's techniques, and this is necessary in order for RSA to be secure. To factor such a number by Shor's quantum algorithm we would need two logical "qubits" for each of the 2048 bits of N. According to estimates by researchers at Microsoft, this would require a circuit with roughly 6 million million (6×10^{12}) physical gates. The exorbitant cost in gates is due to the extreme amount of noise and instability (called "quantum decoherence") in a quantum system, as a result of which a lot of duplication and statistical analysis (called "error correction") is needed in order to get accurate results.

The mathematics of Shor's algorithm is strange and exotic. It's hard enough for us to imagine a space of more than three dimensions. The space where the Shor algorithm computes consists of more than 10^{1200} dimensions. That's more than 10^{1100} times the number of atoms in the Universe! But in the physics of quantum mechanics this is reasonable. If you find this hard to accept, you are not alone. Einstein famously refused to accept the validity of quantum mechanics — despite experimental support for the theory.

The physics and engineering challenges in constructing a quantum computer are far greater than the ones that confronted the developers of the first modern electronic computers. They seem to be of a magnitude much greater than those encountered by the Manhattan Project (the first nuclear bomb) and the Apollo Project (the first humans on the Moon). In the former case, after Einstein's famous letter to U.S. President Roosevelt saying that the U.S. should develop an atom bomb before the Nazis did, less than six years elapsed before the first explosion of an atom bomb at Alamogordo, New Mexico. In the latter case a total of 8 years elapsed between the announcement by U.S. President Kennedy that the country's goal was to travel to the Moon by 1970 and the first landing of a human on the Moon. But in the case of quantum computing already 40 years have elapsed since the possibility was proposed, and over 20 years since intensive work started in hopes of making quantum computing into a reality. Moreover, the eventual outcome is far more uncertain than it was for the Manhattan and Apollo Projects. In fact, some experts believe that quantum decoherence will be an insurmountable obstacle to the construction of a large-scale quantum computer.

Practical Impact of Quantum Computing?

There are two central questions relating to the practicality of quantum computing. First, is there a "killer app"? Here a killer app means an application of sufficient practical value to justify the investment of large amounts of money. The second question is whether a large-scale quantum computer (capable of breaking RSA and ECC) can be constructed in the foreseeable future.

No clear-cut killer app is known. Cryptography is not a killer app, because if quantum computers are ever close to becoming a reality, everyone will switch to new forms of quantum-safe mathematical cryptography. At that point quantum computers will be of no use to hackers and national

intelligence agencies.

There was one other hard computational problem that was frequently cited by promoters of quantum computing as an example of a killer app: the "recommendation problem" of artificial intelligence theory. An example of recommendation systems: My wife loves to read science fiction, and she subscribes to Amazon's "Kindle Unlimited," which provides an unlimited number of e-books for USD 10 per month. Amazon has an algorithm that, using data about her past reading and past reviews and about other readers' reviews, quite accurately can predict what authors and books she will like. Recommendation algorithms have many commercial applications. In 2009 the company Netflix gave a USD 1 million prize to a team that developed an algorithm that improved upon their existing algorithm by 10%.

For a while it was thought that a quantum algorithm could speed up recommendation systems. But in 2018 a brilliant 18-year-old undergraduate student at the University of Texas named Ewin Tang found a non-quantum algorithm that works as well as the quantum algorithm for this problem.

Despite the lack of any "killer apps," the promoters of quantum computation make hugely exaggerated claims about its future uses. First of all, scientists who want to have generous corporate funding of their work often greatly overstate its potential usefulness. In addition, for companies in a capitalist-consumer economy, "hype" (advertising and huge exaggeration) is their standard way of interacting with journalists and with the public.

But even taking into account the expected tendency toward hype, the claims made for quantum computing are extreme.

From IBM: "Quantum computers could one day provide breakthroughs in many disciplines, including materials and drug discovery, the optimization of complex systems, and artificial intelligence."

From Microsoft: "Quantum computing takes a giant leap forward... that will forever alter our economic, industrial, academic, and societal landscape. This has massive implications for research in healthcare, energy, environmental systems, smart materials, and more."

In contrast to the ridiculously exaggerated claims from IBM and Microsoft, a major report in 2018 by the U.S. National Academies of Science, Engineering, and Medicine concluded that "noisy" quantum computers, which do not give reliably correct answers, are much more likely to be feasible in a reasonable timeframe than exact quantum computers with error correction. They warned that unless a "killer app" is found for noisy quan-

tum computations, the interest of the private sector in financing work on quantum computation cannot be sustained. They recommended that if the private companies stop funding that work, then the U.S. government should step in to support quantum computing research, much the way it funds particle physics, pure mathematics, and other fields that are not likely to bring practical benefits in the foreseeable future.

In 2015 a common prediction by promoters of quantum computers was that there is a 50%-50% chance of a large-scale quantum computer with error correction being constructed by 2030. It would cost over USD 10^9 and would require the energy of a dedicated nuclear power plant. Many think that this is far too optimistic, and that maybe mid-century or the 22nd century — or never — would be a more realistic prediction.

How Good Are We at Predicting the Future?

It's commonly assumed that technological change happens faster than expected. The impact of computers on our lives is much greater than anyone would have predicted in the early days of computer technology 60 or 70 years ago. No one then predicted the extreme miniaturization of computers, making computers available to the average person; e-mail (replacing letter-writing); the Internet; or social media. Now in the 21st century when a new technology is discussed we are naturally inclined to think that it will develop rapidly, and that any forward-thinking person or company with money should be investing in the new technology, whether it's cryptocurrency, non-fungible tokens, or quantum computing. However, if we think more carefully about the history of the last 60 or 70 years, we see that technology sometimes progresses much more slowly than predicted.

Science fiction writers spend their lives imagining what a future world will be like, so they should be as good as anyone at predicting future trends. Going back to the 1950s and 1960s, what were they predicting? Let's compare their predictions for the late 20th and early 21st centuries with what has actually happened.

Interplanetary exploration. Prediction: By the year 2000 we would have colonies on the Moon and Mars and perhaps some moons of Jupiter and Saturn and perhaps mining operations on the asteroids. Reality: Between December 1968 and 1972, 24 humans ventured beyond Earth orbit, all through the Apollo Project. Twelve walked on the Moon and twelve flew near the Moon. In the subsequent 50 years no human has gone farther than

earth orbit. In the six decades since Yuri Gagarin orbited the Earth, the most exciting moment in human space exploration occurred in 1969 when the first human set foot on the moon.

Supersonic passenger aircraft. Prediction: By 2000 everyone would travel between North America and Europe by supersonic aircraft. Reality: Supersonic passenger air travel (the French/British "Concorde") existed between 1976 and 2003 (for a small number of rich passengers), and then it ended. It turned out to be neither safe nor economically viable. You might remember the tragic crash of a Concorde in Paris on 25 July 2000.

Nuclear power. Prediction: By 2000 nuclear power will have been scaled down and made inexpensive, so we'd have nuclear-powered cars, ovens, and even watches. Reality: There have been no fundamental advances in nuclear power in a half-century, and the availability of nuclear power in addition to other sources of power has had no impact on daily life.

The Future of Quantum Computing

Will quantum computing see fast progress the way classical computing did? Or will it be more like interplanetary travel and nuclear-powered cars — a nice subject for science fiction, but not a reality any time soon?

In the early years of talk about quantum computing, the accepted measure of progress was going to be the size of the largest numbers that could be factored using Shor's algorithm. After many years and huge sums of money, the progress toward breaking RSA is that the quantum computing people can factor the number $21 = 3 \times 7$. In 2019 an attempt was made to factor $35 = 5 \times 7$, but it failed because of quantum decoherence. (Somewhat larger numbers have been factored by quantum methods, but only by replacing Shor's algorithm with methods that do not scale up to truly large integers.) Not surprisingly, promoters of quantum cryptography have switched to other "benchmarks" of progress related to engineering details. In recent years they have usually avoided talking about progress in actual factoring.

Will We Need New Cryptosystems?

If there is even a small chance that a large-scale quantum computer will be built within 20 or 30 years, we have to worry about any long-term secret whose security depends on public-key cryptography (RSA or ECC).

Most secrets do not have to remain secret for that long. Public-key signatures need to be unforgeable only for a relatively short time. Most national security secrets that require long-term secrecy are protected by private-key cryptography, not RSA or ECC. The currently recommended private-key algorithms, "Secure Hash Algorithm" SHA256 and "Advanced Encryption Standards" AES256, have huge safety margins, and are believed to be safe from any quantum attacks, which, to the best of our knowledge, would not be dramatically faster than non-quantum attacks. For attacks on SHA and AES, classical algorithms have some big advantages: they can be run in parallel, and they can use ASICs (Application-Specific Integrated Circuits). Quantum computers can do neither.

Conclusion

For researchers, the design of a quantum computer and the development of quantum-safe cryptography provide many topics for research in mathematics, physics, computer science, and engineering — a plentiful source of employment for young PhD's. That's a prediction we can make confidently. But will there be a real-world impact? It's far too early to know.